

10/810,164 RA-892

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property
Organization
International Bureau



(43) International Publication Date
10 February 2005 (10.02.2005)

PCT

(10) International Publication Number
WO 2005/013134 A1

(51) International Patent Classification⁷: G06F 12/02
(21) International Application Number: PCT/US2004/023043
(22) International Filing Date: 16 July 2004 (16.07.2004)
(25) Filing Language: English
(26) Publication Language: English
(30) Priority Data: 10/632,474 1 August 2003 (01.08.2003) US

(74) Agents: VINCENT, Lester, J. et al.; Blakely Sokoloff Taylor & Zafman, 12400 Wilshire Boulevard, 7th Floor, Los Angeles, CA 90025 (US).

(71) Applicant (for all designated States except US): INTEL CORPORATION [US/US]; 2200 Mission College Boulevard, Santa Clara, CA 95052 (US).

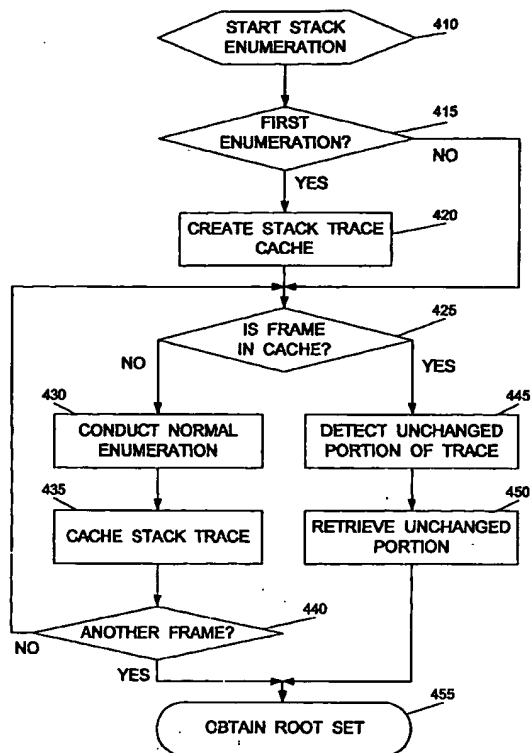
(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.

(72) Inventors: GANSHA, Wu; Room 302, 1# Unite, 18# Building, Liu Fang Nan Li, Chao Yang District, Beijing 100028 (CN). LUEH, Guei-Yuan; 1239 Quail Creek Circle, San Jose, CA 95120 (US).

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI,

[Continued on next page]

(54) Title: METHOD AND APPARATUS FOR IMPROVING THE PERFORMANCE OF GARBAGE COLLECTION USING STACK TRACE CACHE



(57) Abstract: An arrangement is provided for performing root set enumeration in using a stack trace cache when a stack of a thread during garbage collection. During the first root set enumeration in the stack, full stack unwinding may be performed and a stack trace cache may be created to cache stack trace information relating to stack frames. Subsequent sessions of root set enumeration in the stack may access and copy parts or the entire cached stack trace information instead of performing full stack unwinding.

WO 2005/013134 A1



FR, GB, GR, HU, IE, IT, LU, MC, NL, PL, PT, RO, SE, SI,
SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ,
GW, ML, MR, NE, SN, TD, TG).

— *before the expiration of the time limit for amending the
claims and to be republished in the event of receipt of
amendments*

Published:

— *with international search report*

*For two-letter codes and other abbreviations, refer to the "Guid-
ance Notes on Codes and Abbreviations" appearing at the begin-
ning of each regular issue of the PCT Gazette.*

METHOD AND APPARATUS FOR IMPROVING THE PERFORMANCE OF GARBAGE COLLECTION USING STACK TRACE CACHE

RESERVATION OF COPYRIGHT

5

A portion of the disclosure of this patent document contains materials that are subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure, as it appears in the Patent and Trademark Office patent file or
10 records, but otherwise reserves all copyright rights whatsoever.

BACKGROUND

1. FIELD

15

The present invention relates generally to garbage collection and, more specifically, to root set enumeration in a garbage collection process.

2. DESCRIPTION

20

The function of garbage collection, i.e., automatic reclamation of computer storage, is to find data objects that are no longer in use and make their space available for reuse by running programs. Garbage collection is important to avoid unnecessary complications and subtle interactions created by explicit storage allocation, to reduce the complexity of program debugging, and thus to promote fully modular programming and increase software application portability. Because
25 of its importance, garbage collection is becoming an integral part of managed runtime environments.

30

The basic functioning of a garbage collector may comprise three phases. In the first phase, all direct references to objects from currently running programs for all threads may be identified. These references are called roots, or together a root set, and a process of identifying all of such references may be called root set enumeration. In the second phase, all objects reachable from the root set may be searched since these objects may be used in the future. An object that is reachable from any reference in the root set is considered a live object; otherwise

it is considered a garbage object. An object reachable from a live object is also live. The process of finding all live objects reachable from the root set may be referred to as live object tracing (or marking and scanning). In the third phase, storage space of garbage objects may be reclaimed (garbage reclamation). This phase may be conducted either by a garbage collector or a running application (usually called a mutator). In practice, these three phases, especially the last two phases, may be functionally or temporally interleaved and a reclamation technique may be strongly dependent on a live object tracing technique. Depending where root set enumeration occurs, the root set enumeration may be called register root set enumeration (hereinafter register enumeration), heap root set enumeration (hereinafter heap enumeration), or stack root set enumeration (hereinafter stack enumeration). Compared to stack enumeration, overheads incurred by root set enumeration in other storage areas are usually small in a typical application and may be ignored.

When free storage space is running below a limit, garbage collection may be invoked and all threads may be suspended so that root set enumeration for each thread may be started (for concurrent garbage collection, some threads might not be suspended in order to invoke root set enumeration). For stack enumeration for a thread, the stack frame (in the thread's stack) where the thread is suspended becomes a current frame from which stack enumeration may start. All live references in the current frame may be identified and enumerated. After the current frame is enumerated, the next stack frame (i.e., a caller's frame) in a call stack becomes a current frame in which all live references may be identified. This process, which is referred to as stack unwinding, continues until all frames in a call chain are walked through and enumerated.

A stack unwinding mechanism involved in the stack enumeration in a garbage collector unwinds or walks up stack frames of a call stack, one frame at a time, to identify currently active references, i.e., references to form a root set. For some applications, especially those with a large number of threads and a deep call chain per thread, stack unwinding incurs significant runtime overhead for garbage collection. The more threads there are and the deeper the call chain is per thread, the higher the runtime overhead that may be used. Therefore, it is

desirable to improve the efficiency of stack enumeration by reducing the overhead incurred by stack unwinding.

BRIEF DESCRIPTION OF THE DRAWINGS

5

The features and advantages of the present invention will become apparent from the following detailed description of the present invention in which:

10 Figure 1 depicts a high-level framework of an example managed runtime system that uses at least one stack trace cache to improve the performance of garbage collection, according to an embodiment of the present invention;

Figure 2 is an exemplary flow diagram of a high-level process in which a stack trace cache is used during root set enumeration for garbage collection in a managed runtime system, according to an embodiment of the present invention;

15 Figure 3 is a high-level functional block diagram of a stack enumeration mechanism that uses a stack trace cache, according to an embodiment of the present invention;

20 Figure 4 is an exemplary flow diagram of a process in which a stack trace cache is created and used to improve the performance of root set enumeration during garbage collection for a thread, according to an embodiment of the present invention; and

Figures 5(a)-(d) are schematic illustrations of the status of a stack trace cache during different sessions of stack enumeration for a thread, according to an embodiment of the present invention.

25

DETAILED DESCRIPTION

30 An embodiment of the present invention is a method and apparatus for improving the performance of root set enumeration for garbage collection using at least one stack trace cache. The present invention may be used to reduce the overhead of stack enumeration during garbage collection in software applications with a large number of threads and a deep call chain per thread, without much cost. In many software applications, a function call chain in a thread may be

repetitive, that is, caller and callee relationships might not change much, from one garbage collection session to the next. Such a characteristic may result in similar or partly similar stack traces for two consecutive garbage collection sessions. Thus, a stack trace cache may be used for the thread to store stack trace information, which reflects caller-callee relationships in a call chain. The stack trace information may comprise a list of stack frames, context information associated with each frame, current instruction pointer (IP), and/or source line number information. In the first stack enumeration (the first session of garbage collection), full stack unwinding may be performed and a stack trace cache may be created to store trace information for each frame in a stack. In the second or later stack enumeration (the second or later session of garbage collection), part or all work involved in full stack unwinding may be avoided by simply retrieving repetitive portions of stack trace information from the stack trace cache. Also in the second or later stack enumeration, the stack trace cache may be modified to accommodate new traces and/or to update old traces.

Reference in the specification to "one embodiment" or "an embodiment" of the present invention means that a particular feature, structure or characteristic described in connection with the embodiment is included in at least one embodiment of the present invention. Thus, the appearances of the phrase "in one embodiment" appearing in various places throughout the specification are not necessarily all referring to the same embodiment.

Figure 1 depicts a high-level framework of an example managed runtime system that uses at least one stack trace cache to improve the performance of garbage collection, according to an embodiment of the present invention. The managed runtime system 100 may comprise a core virtual machine (VM) 110, at least one Just-In-Time (JIT) compiler 120, a root set enumeration mechanism 130, at least one stack trace cache 140, and a garbage collector 150. The core VM 110 is an abstract computing machine implemented in software on top of a hardware platform and operating system. The use of a VM makes software programs independent from different hardware and operating systems. A VM may be called Java Virtual Machine (JVM) for Java programs, any may be referred to as other names such as, for example, Common Language Infrastructure (CLI) for C# programs. In order to use a VM, a program must first be compiled into an

architecture-neutral distribution format, i.e., intermediate language such as, for example, bytecode for a Java code. The VM interprets the intermediate language and executes the code on a specific computing platform. However, the interpretation by the VM typically imposes an unacceptable performance penalty to the execution of a bytecode because of large runtime overhead processing. A JIT compiler has been designed to improve the VM's performance. The JIT compiler 120 compiles the intermediate language of a given method into a native code of the underlying machine before the method is first called. The native code of the method is stored in memory and any later calls to the method will be handled by this faster native code, instead of by the VM's interpretation.

The root set enumeration mechanism 130 may identify initial live references (roots) in a stack, registers, and other storage areas. The root set enumeration mechanism may be a part of the core VM 110, a part of a garbage collector, a part of both the core VM and the garbage collector, or separate from either the core VM or the garbage collector. The root set enumeration mechanism may comprise a stack enumeration mechanism 135 to identify live references in a stack, a register enumeration component to identify live references in registers, and other components to identify live references in other storage areas. Since stack enumeration may incur more overhead than enumeration in registers and other areas, a stack trace cache 140 may be used to reduce stack enumeration overhead for a thread. A stack trace cache may comprise a storage space in memory. In one embodiment, the stack trace cache may be implemented using dynamic and scalable data structures such as, for example, a linked list.

The stack enumeration mechanism 135 may use a stack trace cache to store stack trace information for a thread. The stack trace information may comprise a list of stack frames, context information associated with each frame, current instruction pointer (IP), and/or source line number information. During the first stack enumeration (for the first garbage collection session), the stack enumeration mechanism may create a stack trace cache and cache the stack trace information in the stack trace cache, while the stack enumeration mechanism walks through and enumerates every frame in a stack, i.e., unwinds the entire stack. During the second or later stack enumeration (for the second or later garbage collection session), the stack enumeration mechanism may first start

to walk through each frame in the stack. For a frame, the stack enumeration mechanism may detect if trace information starting from the frame is cached in the stack trace cache. If the trace information starting from the frame is cached, it may mean that the stack enumeration mechanism can directly use the trace information without further unwinding the stack. On the other hand, if the trace information starting from the frame is not cached, or has changed since the last stack enumeration, the stack enumeration mechanism may modify the stack trace cache to cache the new or updated trace information. After the second or later stack enumeration, the stack enumeration mechanism may update stack trace information in the stack trace cache.

The garbage collector 150 may comprise a live object tracing mechanism 152 and a garbage reclamation mechanism 154. The live object tracing mechanism 152 may search a storage space to find all objects reachable from a root set. Since these objects may be used in the future, they are considered live objects. The live object tracing mechanism may mark these objects as live objects and further scan to find any other live objects these objects can reach. Objects other than live objects are considered garbage objects whose storage space may be reclaimed by the garbage reclamation mechanism 154. In practice, the live object tracing mechanism and the garbage reclamation mechanism may be functionally interleaved and a reclamation technique may be strongly dependent on a live object marking technique. The garbage reclamation mechanism may reclaim garbage objects' space with or without moving live objects to a contingent area at one end of the storage space. In one embodiment, the garbage reclamation mechanism may be a part of a mutator rather than a garbage collector.

Figure 2 is an exemplary flow diagram of a high-level process in which a stack trace cache is used during root set enumeration for garbage collection in a managed runtime system, according to an embodiment of the present invention. At block 210, a garbage collection session may be initiated. At block 220, stack enumeration may be performed using a stack trace cache. A stack trace cache may be created during the first stack enumeration for a thread to cache stack trace information when full stack unwinding is performed. During a later session of stack enumeration for the thread, part or all of cached stack trace information

may be retrieved without conducting full stack unwinding, if it has not changed from the last stack enumeration session. For those frames whose trace information has changed, stack unwinding may be performed and the stack trace cache may be modified to accommodate updated stack trace information. At
5 block 230, root set enumeration in other storage areas may be performed to find live references in other storage areas such as registers. Blocks 220 and 230 may be performed concurrently, or in a reverse order without affecting the spirit of the present invention. At block 240, live references obtained during stack enumeration and root set enumeration in other storage areas may be put together
10 to form a root set. At block 250, a storage space may be searched to mark and scan all live objects reachable from references in the root set. At block 260, all those objects in the storage space that are not marked may be considered garbage and their space may be reclaimed.

Figure 3 is a high-level functional block diagram of a stack enumeration
15 mechanism that uses a stack trace cache, according to an embodiment of the present invention. The stack enumeration mechanism 135 may comprise a stack frame classifier 310, a trace information caching mechanism 320, a stack unwinding mechanism 330, an unchanged trace detecting mechanism 340, and an unchanged trace retrieving mechanism 350. The stack frame classifier 310
20 may use the value of an "in-cache" tag of a stack frame in a stack 370 to decide if the stack frame is newly generated (i.e., not cached) or has already been in a stack trace cache 140 (i.e., cached). The "in-cache" tag may be set up by a compiler when a call is initiated. When the stack frame associated with the call is created for the first time, the compiler may set the value of the "in-cache" tag to be
25 false. When the stack frame classifier classifies a stack frame as "not cached," the stack unwinding mechanism 330 may perform normal stack unwinding for this frame (i.e., enumerate references in this frame and go to the next frame) and save live references in a root set 360. Subsequently, the trace information caching mechanism 320 may cache the trace information of the frame into the stack trace
30 cache and change the value of the "in-cache" tag of the stack frame from false to true, i.e., the stack frame becomes cached. The trace information caching mechanism may comprise a cache creator 325 that creates a stack trace cache when stack enumeration is first conducted. In the second or later session of stack

enumeration, the cache creator may modify the stack trace cache to accommodate newly generated and/or updated stack frames. The stack trace cache may comprise two areas: an identification area to store identifiers of stack frames and a list area to store live references enumerated in each stack frame. In one embodiment, cached stack trace information may be represented by the following data structure:

```

    struct CachedStackFrame {
        <eip, esp> id; /* eip and esp represent IP and SP, respectively */
        Ref_List ref_list; /* cached enumerated references for this frame */
10    } cache[MAX_CALL_DEPTH];

```

where cache[MAX_CALL_DEPTH] array is a thread local object, that is, each thread has its own stack trace cache.

The trace information caching mechanism 320 may also comprise a identification component to identify each stack frame in a stack trace cache with values of a pair of registers, i.e., instruction pointer register (IP) and stack pointer register (SP). The IP may also be referred to as program counter (PC) and may be sufficient to point out source location and method name of a stack frame. The SP points to the current stack position. Since there may be multiple calls (with different stack positions) to a same method on the stack with the same IP (e.g., recursive calls), IP alone might not be able to identify a stack frame and SP may differentiate frames associated with these calls. Thus, using <IP, SP> may be desirable to identify a stack frame uniquely.

When the stack frame classifier 310 classifies a stack frame in the stack 370 as "cached" in the second or a later stack enumeration, the unchanged trace detecting mechanism 340 may search the stack trace cache from top to bottom and detect a starting frame of an unchanged portion of a stack trace. The starting frame of the unchanged portion in the stack trace cache may be detected when values of <IP, SP> of a frame in the stack trace cache is found to be equal to values of <IP, SP> of a current frame. When the starting frame of the unchanged portion is detected in the stack trace cache, the unchanged portion retrieving mechanism 350 may copy references in the unchanged portion directly to the root set 360 without further unwinding the stack.

Figure 4 is an exemplary flow diagram of a process in which a stack trace cache is created and used to improve the performance of root set enumeration during garbage collection for a thread, according to an embodiment of the present invention. At block 410, stack enumeration may be started. At block 415, a
5 decision whether this is the first stack enumeration for the thread may be made. If this is the first stack enumeration for the thread, a stack trace cache may be created for the thread at block 420 before a frame in a stack is evaluated for enumeration starting at block 425; otherwise, frames in the stack may be evaluated (one by one) directly starting at block 425. At block 425, a frame in the
10 stack may be checked to see if the frame is tagged "cached" or "not cached." If the frame is tagged "not cached," normal stack enumeration may be conducted at block 430, where stack unwinding may be started until a "cached" frame is reached in the stack. At block 435, trace information about frames that have been enumerated at block 430 may be cached and the frames may be tagged as
15 "cached." The trace information may comprise an identifier for each frame and enumerated live references in each frame. The identifier of a frame may comprise values of IP and SP. At block 440, the stack is checked to see if there are any frames left to be enumerated. If there are, the next frame in the stack may be evaluated at block 425 and the process between block 425 and block 455 (not
20 including block 455) may be reiterated until no frame is left in the stack. If a frame is found "cached" at block 425, the stack trace cache may be searched to detect a starting point for an unchanged portion of a stack trace at block 445. Once the starting point of the unchanged portion has been detected, references in frames following the starting point in the stack trace cache may be copied to a root set at
25 block 450, without further unwinding the stack. At block 455, references from all frames in the stack may be obtained.

The advantages of using a stack trace cache for stack enumeration may further be illustrated by comparing the process of normal stack enumeration and the process of stack enumeration using a stack trace cache. The former is
30 illustrated by Pseudo Code 1, and the latter is illustrated by Pseudo Code 2.

© 2003 INTEL CORPORATION

```
1  Frame_Context context;  
2  Ref_List allref = {}; /*stack enumeration result list */  
3  ... ..  
4  context.initialize_to_gc_suspend_point();  
5  while (not finished) {  
6      normalize_if_necessary (&context);  
7      /* perform normal stack enumeration */  
8      Ref_List rl = enumerate (&context);  
9      /* copy enumerated references into result list */  
10     allref.append (rl);  
11     finished = unwind_to_next_frame (&context);  
12 }
```

Pseudo Code 1. Normal Stack Enumeration

In Pseudo Code 1, "Frame_Context" on line 1 represents a data structure

of a call frame and contains information such as, for example, saved registers and spilled data. Before stack enumeration for a thread starts, "context" is initialized to a frame where the thread is suspended through "initialize_to_gc_suspend_point()" in line 4. "normalize_if_necessary()" in line 6 normalizes a frame's SP from its
 5 suspended position to its base position, if necessary. Enumeration of references in this frame may start from normalized SP position. At the end of enumeration for this frame, "unwind_to_next_frame()" unwinds the stack to a caller's frame in a

© 2003 INTEL CORPORATION

```

1  Frame_Context context;
2  Ref_List allref = {}; /* initialize stack enumeration result list */
3  /* temporary cache to save stack frames not in cache */
4  CachedStackFrame notcached[MAX_CALL_DEPTH];
5  ... ..
6  context.initialize_to_gc_suspend_point();
7  while (not finished) {
8      if (context.in_cache == false) {
9          normalize_if_necessary (&context);
10         /* perform a normal enumeration */
11         Ref_List rl = enumerate (&context);
12         /* copy enumerated references into result list */
13         allref.append (rl);
14         /* add information of this frame to temporary cache */
15         add_to_cache (notcached, <context.eip, context.esp>, rl);
16         finished = unwind_to_next_java_frame (&context);
17     } else {
18         /* find the starting point of reusable trace in cache */
19         for (j = top; j >= 0; j--)
20             if (context.<eip, esp> == cache[j].<eip, esp>) break;
21         /* copy enumerated references in cache to result list */
22         for (; j >= 0; j--)
23             allref.append (cache[j].ref_list);
24         /* stop unwinding, and jump out of loop */
25         break;
26     }
27 }
28 /* update the cache with information in temporary cache */
29 update_cache (cache, notcached);

```

Pseudo Code 2. Stack Enumeration Using Stack Trace Cache

call chain and retrieve the caller's context. The enumeration process continues

until all frames in the call chain are enumerated. In other words, Pseudo Code 1 illustrates a process of full stack unwinding.

Pseudo Code 2 differs from Pseudo Code 1 by utilizing a stack trace cache (shown in lines 18-25) to increase the opportunity of avoiding part or all of stack unwinding. Therefore, overheads incurred by stack enumeration using a stack trace cache may be smaller than overheads incurred by stack enumeration through full stack unwinding.

Figures 5(a)-(d) are schematic illustrations of the status of a stack trace cache in different sessions of stack enumeration for a thread, according to an embodiment of the present invention. Figure 5(a) illustrates a stack before the first session of stack enumeration is started. When stack enumeration is invoked for the first time for a thread, all frames in the stack are tagged "not cached" by a compiler. Thus, full stack unwinding is performed during the first stack enumeration, after which all frames are tagged with "cached" because a stack trace cache has been created and enumerated references for each frame has been cached in the stack trace cache, as shown in Figure 5(b). Each frame in the stack trace cache has an identification (id) area to store frame ids (i.e., values of IP and SP), and a list area to store enumerated references for each frame. Figure 5(c) illustrates status of each frame in the stack and the stack trace cache in a later session of stack enumeration. When this later session of stack enumeration is initiated, it may be not necessary to perform full stack unwinding because some frames (e.g., 4 frames corresponding to methods T.main, A.a, B.b, and C.c in Figure 5(c)) remain unchanged in the stack and may be directly retrieved from the stack trace cache without conducting stack unwinding. Only information for another three frames (i.e., frames corresponding to methods G.g, H.h, and I.i) are not cached. Thus only partial stack unwinding (up to the frame corresponding to method B.b) is needed. Additionally, the stack trace cache may be modified to cache information on frames being enumerated during the partial stack enumeration. Figure 5(d) illustrates the status of the stack and the stack trace cache after this later session of stack enumeration.

Although the present invention is concerned with using stack trace caches for root set enumeration in a stack during garbage collection, persons of ordinary skill in the art will readily appreciate that the present invention may be used for

reducing overheads incurred by any process involving stack unwinding such as, for example, exception handling, caller-callee relationship detecting, etc. Additionally, the present invention may be used for automatic garbage collection in any systems such as, for example, managed runtime environments running
5 Java, C#, and/or any other programming languages.

Although an example embodiment of the present invention is described with reference to block and flow diagrams in Figures 1-5(d) and Pseudo Codes 1-2, persons of ordinary skill in the art will readily appreciate that many other methods of implementing the present invention may alternatively be used. For
10 example, the order of execution of the blocks in flow diagrams or steps in pseudo codes may be changed, and/or some of the blocks in block/flow diagrams and the steps in pseudo codes described may be changed, eliminated, or combined.

In the preceding description, various aspects of the present invention have been described. For purposes of explanation, specific numbers, systems and
15 configurations were set forth in order to provide a thorough understanding of the present invention. However, it is apparent to one skilled in the art having the benefit of this disclosure that the present invention may be practiced without the specific details. In other instances, well-known features, components, or modules were omitted, simplified, combined, or split in order not to obscure the present
20 invention.

Embodiments of the present invention may be implemented on any computing platform, which comprises hardware and operating systems. The hardware may comprise a processor, a memory, a bus, and an I/O hub to peripherals. The processor may run a compiler to compile any software to the
25 processor-specific instructions. Processing required by the embodiments may be performed by a general-purpose computer alone or in connection with a special purpose computer. Such processing may be performed by a single platform or by a distributed processing platform. In addition, such processing and functionality can be implemented in the form of special purpose hardware or in the form of
30 software.

If embodiments of the present invention are implemented in software, the software may be stored on a storage media or device (e.g., hard disk drive, floppy disk drive, read only memory (ROM), CD-ROM device, flash memory device,

digital versatile disk (DVD), or other storage device) readable by a general or special purpose programmable processing system, for configuring and operating the processing system when the storage media or device is read by the processing system to perform the procedures described herein. Embodiments of

5 the invention may also be considered to be implemented as a machine-readable storage medium, configured for use with a processing system, where the storage medium so configured causes the processing system to operate in a specific and predefined manner to perform the functions described herein.

While this invention has been described with reference to illustrative

10 embodiments, this description is not intended to be construed in a limiting sense. Various modifications of the illustrative embodiments, as well as other embodiments of the invention, which are apparent to persons skilled in the art to which the invention pertains are deemed to lie within the spirit and scope of the invention.

15

CLAIMS

What is claimed is:

- 5 1. A method for automatically collecting garbage objects, comprising:
performing root set enumeration in a stack of a thread using a stack trace
cache;
obtaining a root set by combining references obtained from root set
enumeration for all threads;
tracing live objects that are reachable from references in the root set; and
10 reclaiming storage space occupied by objects other than the live objects.
2. The method of claim 1, wherein the root set comprises currently live
references in stacks, registers, and other storage areas for all threads.
- 15 3. The method of claim 1, wherein performing root set enumeration in a
stack of a thread using a stack trace cache comprises:
determining if a stack frame is cached in the stack trace cache;
detecting a portion of stack trace information that has not changed since
the last stack enumeration in the stack trace cache and retrieving the unchanged
20 portion from the stack trace cache, if the stack frame is cached in the stack trace
cache; and
performing normal root set enumeration in the stack, if the stack frame is
not cached in the stack trace cache.
- 25 4. The method of claim 1, wherein tracing the live objects comprises
marking the live objects and further scanning objects that are reachable from the
marked objects.
5. A method for performing root set enumeration in a stack of a thread for
30 automatic garbage collection, comprising:
initializing a stack enumeration list;
creating a stack trace cache for the thread; and

performing root set enumeration in the stack using the stack trace cache to produce an updated stack enumeration list.

5 6. The method of claim 5, further comprising tagging a stack frame to indicate whether the stack frame has been cached in the stack trace cache or is newly generated in the stack and has not been cached in the stack trace cache.

7. The method of claim 5, wherein the stack enumeration list comprises currently live references in the stack.

10

8. The method of claim 5, wherein creating the stack trace cache comprises:

determining if the root set enumeration in the stack is performed for the first time;

15 dynamically creating the stack trace cache and caching information of stack frames in the stack trace cache, if the root set enumeration in the stack is performed for the first time; and otherwise, reusing the stack trace cache for the root set enumeration in the stack.

20 9. The method of claim 8, wherein caching information of stack frames comprises adding an identifier to each frame using values of an instruction pointer register and a stack pointer register.

25 10. The method of claim 8, wherein reusing the stack trace cache comprises at least one of the following: using the stack trace cache without a change, and modifying the stack trace cache to accommodate new and updated stack frames.

30 11. The method of claim 5, wherein performing root set enumeration in the stack using the stack trace cache comprises:
determining if a stack frame is cached in the stack trace cache based on a tag of the stack frame;

copying an unchanged portion of stack trace information from the stack trace cache to the stack enumeration list, if the value of the tag indicates that the stack frame is cached in the stack trace cache; and otherwise performing normal root set enumeration in stack.

5

12. The method of claim 11, wherein copying the unchanged portion of the stack trace information comprises detecting a portion of the stack trace information that has not changed since the last stack enumeration in the stack trace cache and retrieving the unchanged portion from the stack trace cache.

10

13. The method of claim 11, wherein performing normal stack enumeration comprises:

unwinding the stack until reaching a frame that is cached in the stack trace cache;

15

caching information of unwound frames in the stack trace cache; and adding enumerated references from the unwound frames into the stack enumeration list.

20

14. A managed runtime system, comprising:

a virtual machine;

a just-in-time compiler;

a root set enumeration mechanism capable of obtaining a root set;

at least one stack trace cache capable of storing trace information of stack frames when the stack frames are enumerated during root set enumeration; and

25

a garbage collector capable of tracing live objects reachable from the root set and reclaiming storage space occupied by objects other than the live objects.

30

15. The system of claim 14, wherein the root set enumeration mechanism comprises a stack enumeration mechanism capable of performing root set enumeration in stack.

16. The system of claim 14, wherein the stack trace cache comprises an identification area to store stack frame identifiers and a list area to store enumerated references for stack frames.

5 17. A system for root set enumeration in a stack during garbage collection, comprising:

 a stack frame classifier capable of classifying a stack frame based on a tag of the stack frame;

 a stack information caching mechanism capable of caching information of
10 the stack frame in a stack trace cache, if the value of the tag indicates that the stack frame is not cached;

 an unchanged trace detecting mechanism capable of detecting a portion of stack trace information that has not changed since the last stack enumeration in the stack trace cache, if the value of the tag indicates that the stack frame is
15 cached; and

 an unchanged trace retrieving mechanism capable of retrieving the unchanged portion of stack trace information from the stack trace cache.

20 18. The system of claim 17, further comprising a stack unwinding mechanism capable of unwinding a stack until reaching a stack frame that is cached in the stack trace cache.

 19. The system of claim 17, wherein the trace information caching mechanism comprises:

25 a cache creator capable of creating the stack trace cache when root set enumeration in the stack is performed for the first time; and

 a frame identification component capable of identifying a frame in the stack trace cache.

30 20. An article comprising: a machine accessible medium having content stored thereon, wherein when the content is accessed by a processor, the content provides for automatically collecting garbage objects by:

performing root set enumeration in a stack of a thread using a stack trace cache;

obtaining a root set by combining references obtained from root set enumeration for all threads;

5 tracing live objects that are reachable from references in the root set; and
reclaiming storage space occupied by objects other than the live objects.

21. The article of claim 20, wherein the root set comprises currently live references in stacks, registers, and other storage areas for all threads.

10

22. The article of claim 20, wherein content for performing root set enumeration in a stack of a thread using a stack trace cache comprises content for:

determining if a stack frame is cached in the stack trace cache;

15 detecting a portion of stack trace information that has not changed since
the last stack enumeration in the stack trace cache and retrieving the unchanged
portion from the stack trace cache, if the stack frame is cached in the stack trace
cache; and

performing normal root set enumeration in the stack, if the stack frame is
20 not cached in the stack trace cache.

23. The article of claim 20, wherein content for tracing the live objects
comprises content for marking the live objects and further scanning objects that
are reachable from the marked objects.

25

24. An article comprising: a machine accessible medium having content
stored thereon, wherein when the content is accessed by a processor, the content
provides for performing root set enumeration in a stack of a thread for automatic
garbage collection by:

30 initializing a stack enumeration list;

creating a stack trace cache for the thread; and

performing root set enumeration in the stack using the stack trace cache to
produce an updated stack enumeration list.

25. The article of claim 24, further comprising content for tagging a stack frame to indicate whether the stack frame has been cached in the stack trace cache or is newly generated in the stack and has not been cached in the stack trace cache.

26. The article of claim 24, wherein the stack enumeration list comprises currently live references in the stack.

27. The article of claim 24, wherein content for creating the stack trace cache comprises content for:
determining if the root set enumeration in the stack is performed for the first time;
dynamically creating the stack trace cache and caching information of stack frames in the stack trace cache, if the root set enumeration in the stack is performed for the first time; and otherwise,
reusing the stack trace cache for the root set enumeration in the stack.

28. The article of claim 27, wherein content for caching information of stack frames comprises content for adding an identifier to each frame using values of an instruction pointer register and a stack pointer register.

29. The article of claim 27, wherein content for reusing the stack trace cache comprises at least one of the following: content for using the stack trace cache without a change, and content for modifying the stack trace cache to accommodate new and updated stack frames.

30. The article of claim 24, wherein content for performing root set enumeration in the stack using the stack trace cache comprises content for:
determining if a stack frame is cached in the stack trace cache based on a tag of the stack frame;

copying an unchanged portion of stack trace information from the stack trace cache to the stack enumeration list, if the value of the tag indicates that the stack frame is cached in the stack trace cache; and otherwise performing normal root set enumeration in stack.

5

31. The article of claim 30, wherein content for copying the unchanged portion of the stack trace information comprises content for detecting a portion of the stack trace information that has not changed since the last stack enumeration in the stack trace cache and content for retrieving the unchanged portion from the stack trace cache.

10

32. The article of claim 30, wherein content for performing normal stack enumeration comprises content for:

unwinding the stack until reaching a frame that is cached in the stack trace cache;

15

caching information of unwound frames in the stack trace cache; and adding enumerated references from the unwound frames into the stack enumeration list.

1/8




FIGURE 1




FIGURE 2

3/8




FIGURE 3

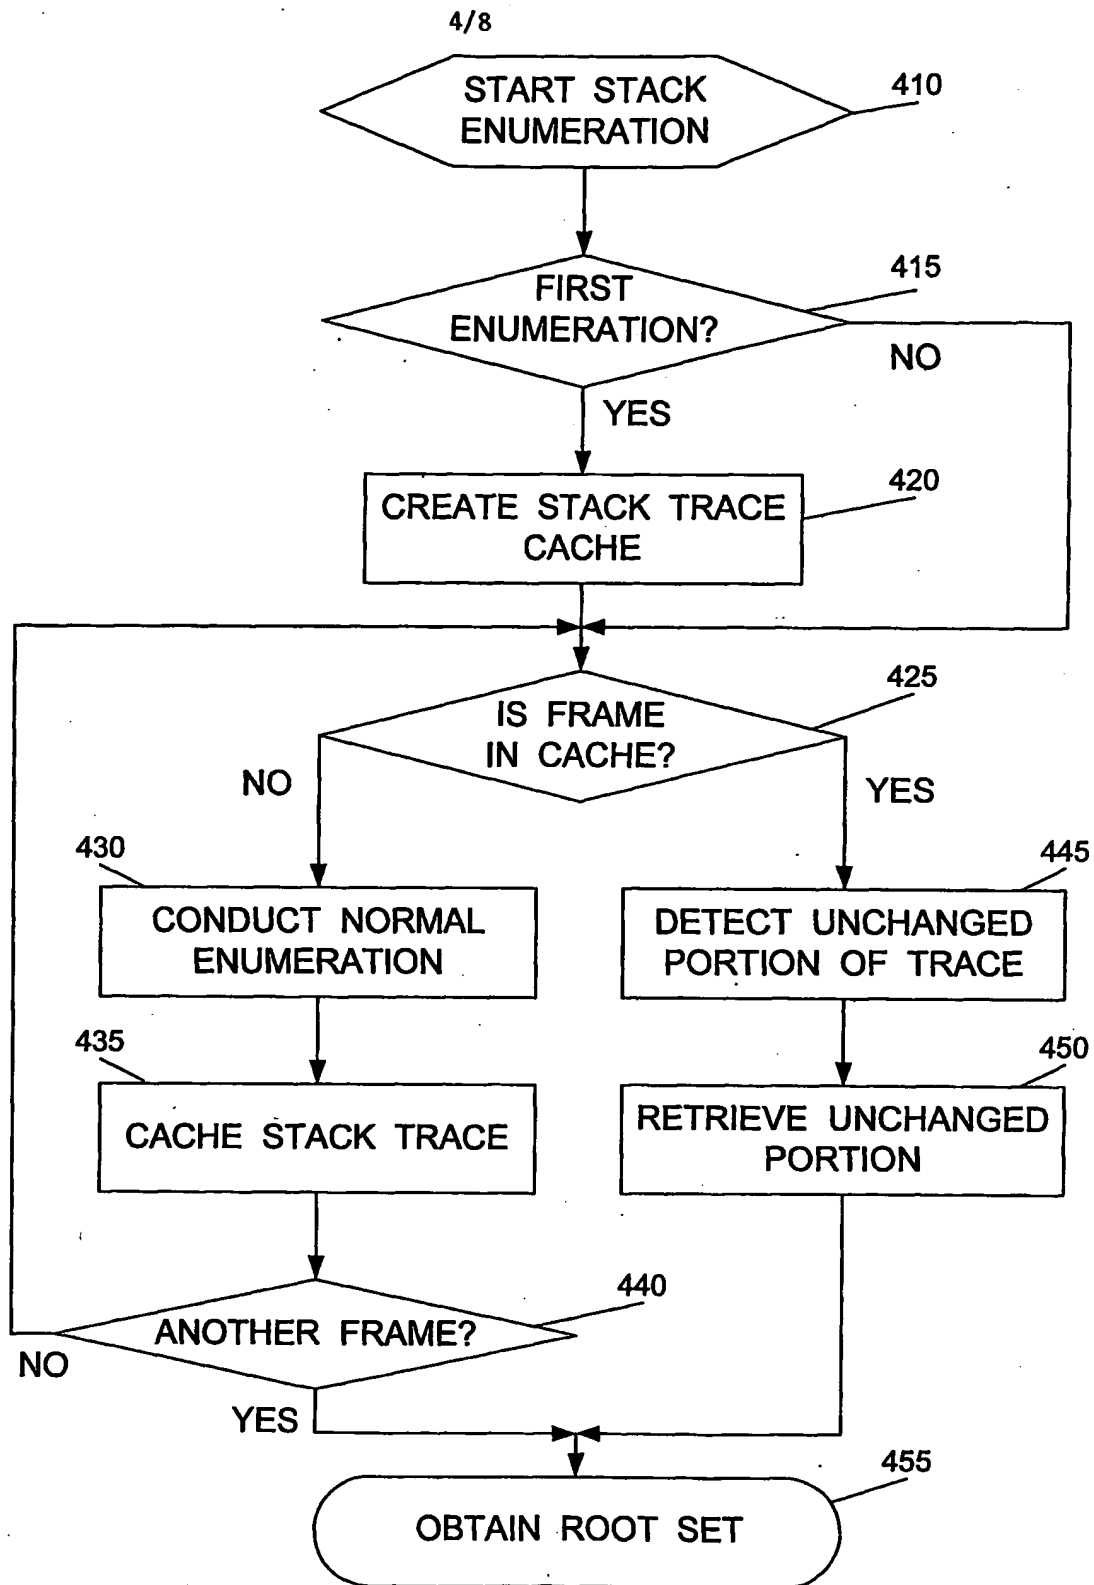



FIGURE 4

5/8



6/8




FIGURE 5(b)




FIGURE 5(c)

8/8




FIGURE 5(d)

INTERNATIONAL SEARCH REPORT

International Application No
PCT/US2004/023043

A. CLASSIFICATION OF SUBJECT MATTER

IPC 7 606F12/02

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 606F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the International search (name of data base and, where practical, search terms used)

EPO-Internal, PAJ

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 6 101 580 A (TREMBLAY MARC ET AL) 8 August 2000 (2000-08-08) figures 2,4 column 1, line 41 - column 1, line 44 column 2, line 13 - column 2, line 27 column 2, line 35 - column 2, line 39 column 4, line 11 - column 4, line 18 column 5, line 30 - column 5, line 32 column 5, line 44 - column 5, line 52 column 10, line 8 - column 15, line 22	1-32
X	US 2003/126352 A1 (BARRETT KIM A) 3 July 2003 (2003-07-03) paragraph '0051!	1,5,14, 17,20,24
A	US 6 470 361 B1 (ATTANASIO CLEMENT RICHARD ET AL) 22 October 2002 (2002-10-22) the whole document	1-32

☐ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents:

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the International filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the International filing date but later than the priority date claimed

T later document published after the International filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

X document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

Y document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

Z document member of the same patent family

Date of the actual completion of the International search

30 November 2004

Date of mailing of the International search report

06/12/2004

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax (+31-70) 340-3016

Authorized officer

Filip, L

INTERNATIONAL SEARCH REPORT

International Application No
PCT/US2004/023043

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 6101580	A	08-08-2000	DE 69802056 D1 22-11-2001 EP 0912942 A1 06-05-1999 JP 2000513854 T 17-10-2000 WO 9848354 A1 29-10-1998
US 2003126352	A1	03-07-2003	NONE
US 6470361	B1	22-10-2002	NONE